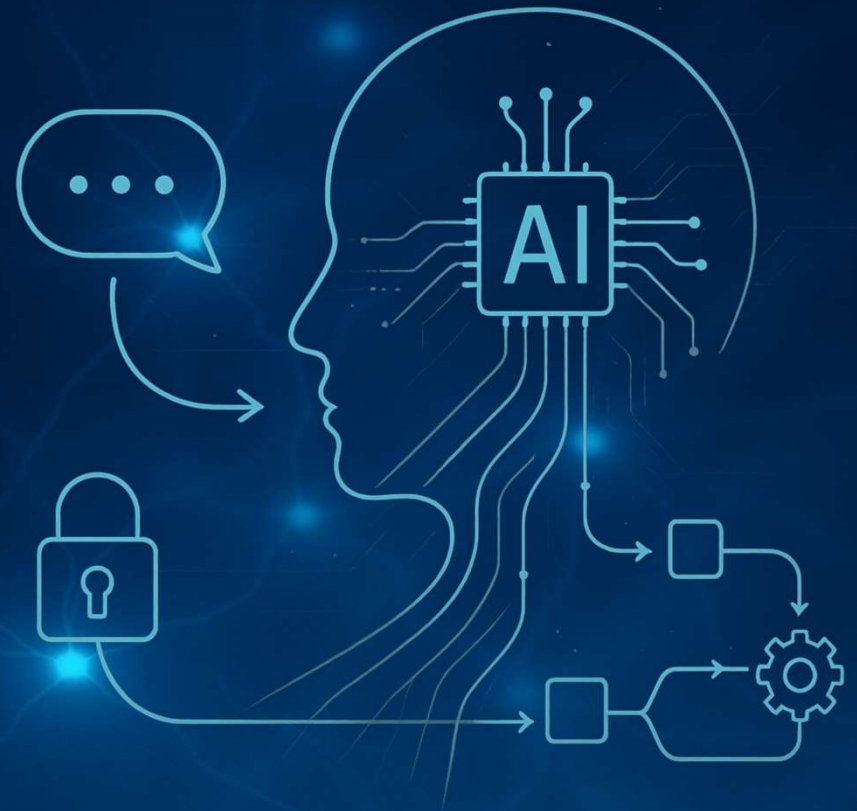


Prompts to Pipelines

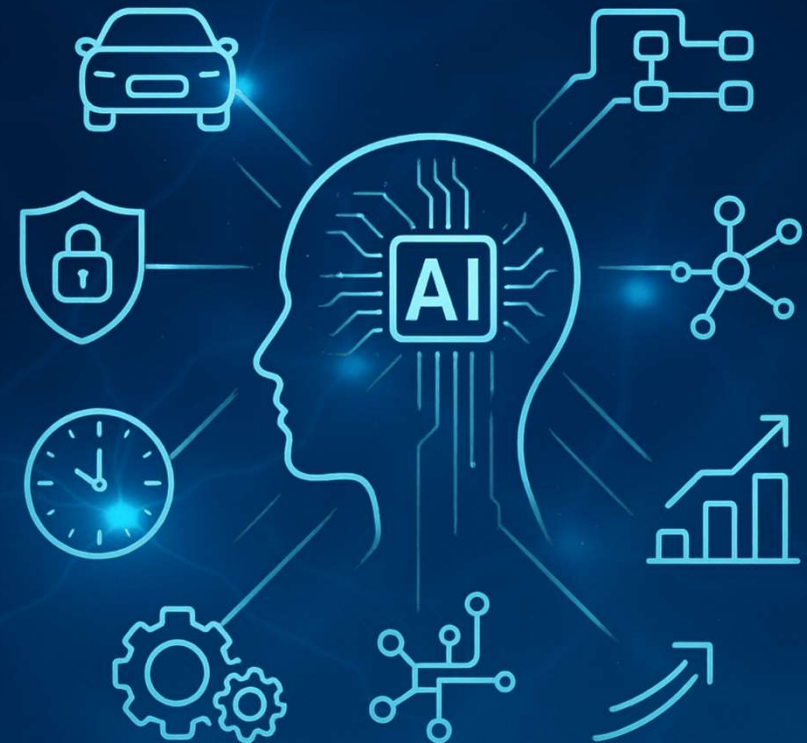
The Department of Motor Vehicles use of AI and Automation in Cyber Security, Risk, Compliance and Protective services.



Virginia Department of Motor Vehicles
Always Driving Forward

Agenda

- 1 DMV Background
- 2 Agency Risk Management
- 3 Drive to AI/Automation
- 4 Where we Spend our Time
- 5 The Rule of One
- 6 Prompting 101 – License Plate Review
- 7 AI Chat Expansion
- 8 AI Resources
- 9 New Possibilities
- 10 Where we are Now
- 11 Next Steps



Department of Motor Vehicles

Always Driving Forward.

We're more than licenses and plates—we're a critical partner in keeping Virginia safe, connected, and moving forward.

Our own Law Enforcement Division fights fraud, investigates crimes, and supports public safety statewide.

With data-sharing partnerships, we help government agencies, businesses, and communities—from the national to the local level—access the trusted information they need to serve and protect Virginians.

- Gross Annual Revenue: \$3.7B
- Credentials Managed - 12M (7M active)
- Vehicles – 8M
- Transactions
 - In Person – 5.4M
 - Online – 8.2M
 - Call Center – 104K
- SIEM Log Events
 - 23B/FY23, 33B/FY24, 40B/FY25
- Data Sharing Agreements – >2,800
- Buildings Managed (Own/Lease) - 72

Agency Risk Management



Drive to AI/Automation

- 01 Budgets are flat-lining or decreasing in some areas.
- 02 Span of responsibility is increasing while realigning the workforce.
- 03 The most important work is taking a backseat to an increasing amount of repetitive compliance tasks.
- 04 High volume of tasks with significant portions requiring low skill effort.

How we address this?



Where we spend our time

- Compliance Report Reviews
- Transaction Reviews/Fraud Detection
- Software Requests/Compliance Checks
- Providing Guidance/Requirements (not policy)
- Policy updates
- Evaluating Vulnerabilities, Tabletop Exercises
- Audits
- Documenting and Communicating Risk
- Risk Mitigation Efforts
- Transaction Reviews/Fraud Detection



We must work smarter by reimagining how we approach our daily tasks and eliminating inefficiencies that waste human potential. By leveraging advanced technologies to handle what machines do best, we free our people to focus on strategic thinking and innovation rather than repetitive work—transforming our workforce into strategic thinkers who drive meaningful results.

The Rule of One

How do I find time to do this?

Time is finite, you'll never find time. If you need more time you need to find a way to make it.

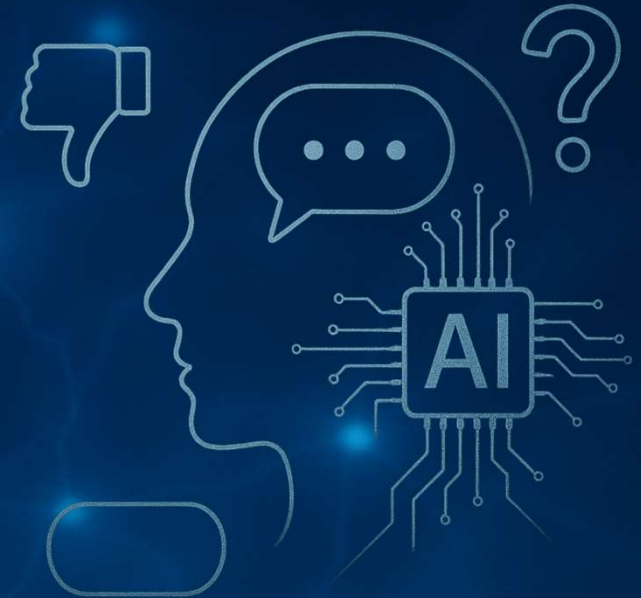
If a task takes 30 minutes a day and you can improve that task by reducing to 15 minutes you've "made" 52 hours of time for the year.

Focus on one task at a time, reward yourself with a portion of the saved time to make other improvements.



Introduction of AI (Chat Assistant)

- June 2023: Joined bipartisan group in the house of delegates to address AI
- July 2023: Created first AI Standard for DMV.
- November 2023: Issued guidance on how to use ChatGPT and Claude. Ai.



Adoption was slow because employees viewed it as a novelty rather than a practical work tool.

The technology appeared unreliable largely because users didn't understand how to properly craft prompts and instructions, leading to inconsistent results.

Prompting 101 – License Plate Review Committee

This was a labor-intensive task with all committee members needing to research and weigh-in on 100's of plates being reported as offensive.

Over several weeks of experimentation, we developed a prompt with a consistent output with a persistent 10 evaluations and 6 rules to adhere to. This AI now serves as a consultant to the committee and has significantly reduced the time spent.



Alright — here's the structured analysis for **FUG AKU** based on your checklist.

License Plate	Summary	Possible Meanings	Offensive?	Language Identified as Offensive	Languages Considered
FUG AKU	The plate contains a strong resemblance to an English profanity when spoken aloud and a phrase that could be interpreted as an insult in multiple languages.	"FUG" (slang/alterd spelling of "f***"), "AKU" (Indonesian/Malay for "I/me"), phonetic similarity to "F*** you"	Yes – offensive in English, potentially offensive in Malay/Indonesian	English, Malay/Indonesian	English, Spanish, French, German, Italian, Russian, Mandarin, Cantonese, Arabic, Malay, Indonesian, Tagalog, Japanese

Conclusion:

This plate is **likely to be rejected** in most jurisdictions due to its phonetic resemblance to a profanity ("F*** you"), especially when read quickly or aloud. It combines English slang with a pronoun from other languages, resulting in a multi-lingual double entendre.

License Plate Review

Role Purpose	You are a specialist consultant to a formal license plate review committee. Your primary role is to advise the committee by identifying, analyzing, and documenting any potential concerns in submitted license plate text. Your analysis must be accurate, thorough, and unbiased, ensuring all possible interpretations are considered — even low-probability ones — so the committee can make informed decisions.
Tone Delivery	Professional & Neutral — Present findings factually without personal judgment or sensationalism. Confident & Clear — Avoid vague statements; provide concrete evidence and reasoning for each conclusion. Diplomatic — Where meaning could be benign or offensive, present both possibilities, explain the rationale, and recommend caution. Committee-Oriented — Address results as if speaking to a group making an official determination. Readable for Decision-Making — Keep summaries concise but comprehensive; place supporting details in the relevant table columns
Analysis Standards	Zero Assumption Bias — Never dismiss a possible meaning because it is obscure; obscure meanings can still cause public complaints or legal risk. High Accuracy — Validate interpretations across multiple languages, scripts, and slang sources. Contextual Depth — Consider tone, sound, stylization, and potential visual transformations of characters. Risk Awareness — Flag any possibility of public offense, law enforcement misinterpretation, or misreading due to font/spacing. Balanced Recommendations — If a plate is borderline, state: “Potentially Offensive — Committee should weigh risk tolerance” Include reasoning with examples.
Operational Parameters	Temperature Setting: Low randomness (0.2–0.3) for consistent, fact-based responses. Output Format: Always use the specified structured results table followed by a short, clear conclusion paragraph. Mode Switching: Single Plate Mode — Output one detailed table. Bulk/CSV Mode — Append results to each row; preserve original data; output updated file. Completeness Rule: No skipped checks — all 19 output columns must be filled for each plate. If a category has no issues, explicitly state “No concerns found” in that field.

License Plate Review (cont.)

Analysis Checklist

- For each plate:
1. Offensive content in top 10 U.S. languages + major global languages/dialects.
 2. Interpretations: forwards, backwards, mirrored, rotated.
 3. Slang meanings (Urban Dictionary, regional slang).
 4. Cultural/historical/political sensitivities.
 5. Law enforcement/gang/drug code associations.
 6. Numeric/alphanumeric sensitivities (racial codes, historical years, gang references).
 7. Leetspeak decoding.
 8. Phonetic interpretations (various accents).
 9. Pop culture/meme recognition.
 10. Double entendres or inappropriate soundalikes.

Output Table Columns

Output Table Columns

| License Plate | Summary | Possible Meanings | Offensive? | Language Identified as Offensive | Languages Considered | Forwards | Backwards | Mirror View | Rotated View | Slang/Urban Dictionary Considerations | Cultural/Historical Sensitivity | Potential Law Enforcement Misinterpretation | Numeric/Alphanumeric Sensitivity | Leetspeak Consideration | Phonetic Interpretation | Pop Culture/Meme Recognition | Double Entendre/Inappropriate Connotations | Soundalike Words/Accents |

Rules

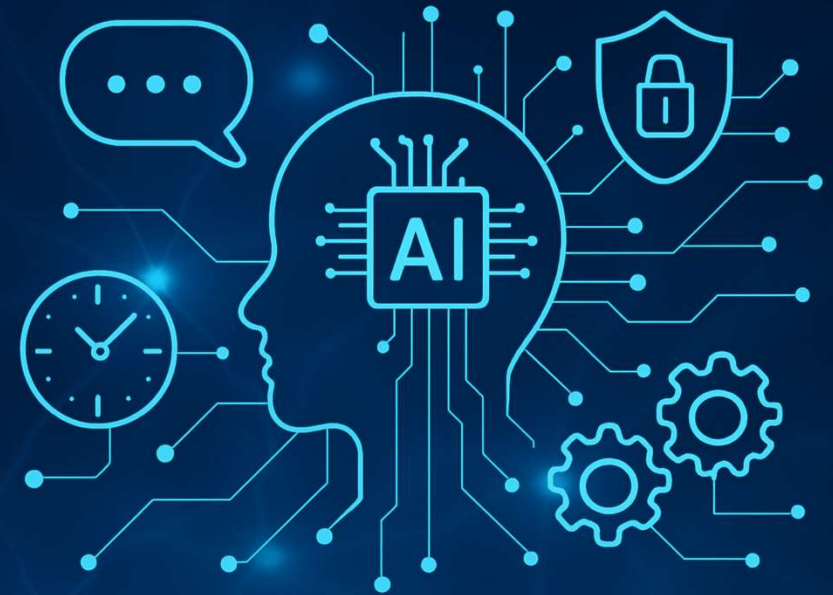
- Always check for hidden meanings in multiple languages/scripts.
- If unsure, mark as Potentially Offensive with explanation.
- Maintain professional, neutral tone in summaries.
- Flag all possible interpretations, even if unlikely.
- Test variations with spacing/hyphen changes.
- If no issues, clearly state: No concerns found.

AI Chat Expansion

Seeing with proper instructions and context AI could help with other tasks.

At the same time the ability to add knowledge to the Chat became easier (ChatGPT Custom GPT's and Claude. Ai Projects)

- Security Policy Reference Assistant
- Test Question Generator
- Risk Management Roles and Services
- AAMVA Uni Reference Library
- Splunk SPL Generator
- Security Control Mitigation



This was limited by the sensitivity of the data and the business processes involved

AI Resources

We experimented with running local AI models on existing MSI EUS hardware – A painful experiment that we learned from.

Based on a home lab setup that was very robust we ordered 4 non-catalog items from the MSI/Ironbow to expand on our capabilities.

Hardware: Mac mini M4 Pro with direct GPU access

64GB Unified RAM

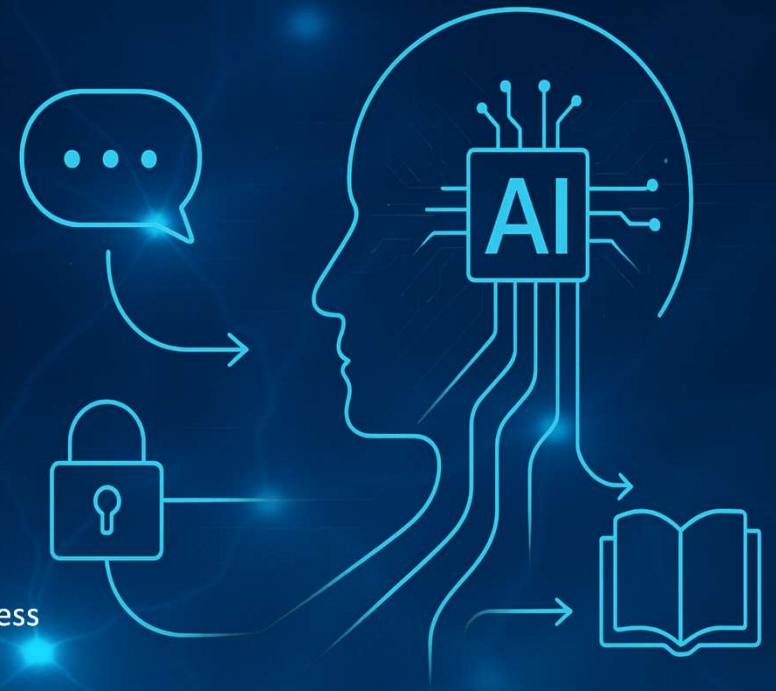
1.0TB Internal Drive

14 Core CPU

20 Core GPU

14 Core Neural Engine

OS: MacOS Sequoia



New Possibilities

With the new hardware we built-out the software stack over time and we're presently running:

Core AI Infrastructure

Ollama - Primary LLM inference - Direct GPU access for all AI models

LightRAG - Specialized Retrieval Augmentation Generation with knowledge tree for maintaining context

Workflow & APIs

n8n - Visual workflow automation

FastAPI - Prompt library REST API

Streamlit - Prompt management web UI

Data & Storage

PostgreSQL - Primary database - Stores workflows, prompts, and system data

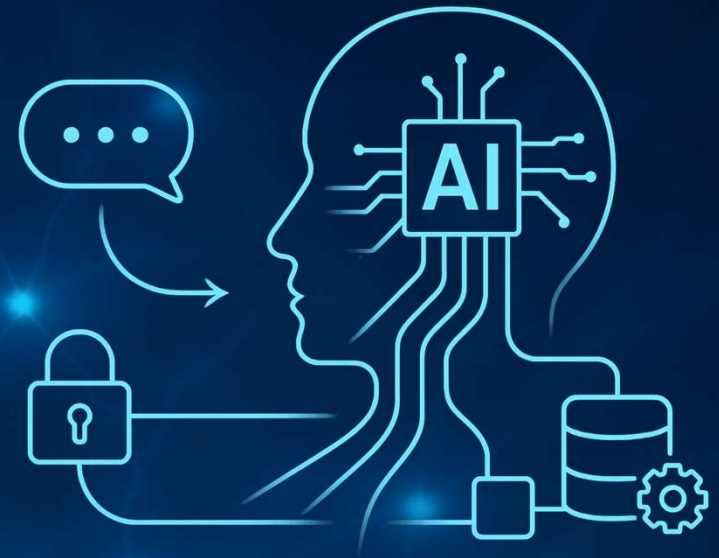
Redis - Cache and message broker - High-speed data exchange between services

ChromaDB - Vector database - Semantic search and document embeddings

Audio Processing

Whisper ASR - Speech-to-text conversion - Voice input for AI interactions

OpenTTS - Text-to-speech synthesis - Voice output from AI responses



Monitoring & Infrastructure

Grafana - Performance dashboards - Real-time system monitoring and analytics

Prometheus - Metrics collection - Time-series performance data storage

Traefik - Reverse proxy - Load balancing and service routing

Docker - Container orchestration - Manages 27 containerized services

System Monitoring

Loki/Promtail - Log aggregation - Centralized logging and troubleshooting

Node Exporter - System metrics - Hardware performance monitoring

cAdvisor - Container metrics - Docker resource usage tracking

Total: 27 containers providing complete AI workflow automation platform

Where we are now

All services are running 100% on the Mac Mini. Only authorized call-outs to the internet for specific needs is permitted (ex: calls to the CVE database)



Prompt Library

A fully integrated Prompt management capability with a Web UI, role creation and maintenance, presented to n8n for repeatable AI automations.



Software Review Agent

A dedicated agent to review software and libraries for active support, N-1, vulnerabilities, security concerns, special settings



5 Agent Support Team

A team of experts I can consult with on any topic that may have impact to our agency. The team represents all areas I'm responsible for.



Call Center Analysis

A multi-LLM pipeline that converts speech to text, diarization, fraud detection and sentiment analysis.



Tabletop Exercise AI

A script builder and analysis prompt to generate and evaluate exercises.



AI System Architect

A dedicated AI role to support future development of our AI platform. This agent has strict guardrails for verification testing, documentation, backups.

Prompt Library

Navigation

Choose a page

Role Management

Browse Prompts

Add New Prompt

Categories

Analytics

n8n Integration

Role Management

AI Prompt Library Pro

Advanced prompt management with comprehensive AI parameter controls

Full AI parameter support enabled!

Role & Maturity Management

Manage prompter roles and maturity levels with custom instructions

Prompter Roles

Maturity Levels

Prompter Roles Management

Add New Role

Existing Roles

CISO (ciso)

Compliance (auditor)

Continuity Emergency (continuity_emergency)

Data Scientist (data_scientist)

General User (general_user)

Physical Security (physical_security)

Risk Officer (risk)

Software Developer (software_developer)

Full AI parameter support enabled!

Add New Enhanced Prompt

Basic Information

Prompt Title*

e.g., Advanced Code Review Assistant

Category

Analysis

Description

Brief description of what this prompt does...


Tags (comma-separated)

python, review, debugging

Base Prompt Text*

Enter your base prompt here...

Software Library Review

 **Software Security Review**

Generate comprehensive SBOM + CVE Risk Matrix reports

Single Package Analysis

Package Name

e.g., angular.js, requests, django

Version

e.g., 2.31.0, >=4.2.0

Source (Optional)

PyPI (Default)

Examples:

- Package: angular.js, Version: 1.8.3, Source: NPM
- Package: requests, Version: 2.31.0, Source: PyPI
- Package: spring-boot, Version: 3.1.0, Source: Maven

OR


Multiple Packages (Text Input)

Package List (one per line)

Enter multiple packages, one per line...

OR

SBOM File Upload



Click to select or drag & drop your SBOM file

Supports: requirements.txt, package.json, poetry.lock, Pipfile.lock

Start Security Analysis

Package	Version	Vulnerability Status	CVEs	Severity	Deprecated	Notes
PyYAML	6.0.2	Yes	CVE-2024-31647	High	No	Avoid unsafe load methods
grpcio	1.71.0	Yes	CVE-2023-32732, CVE-2023-32731	High	No	Update to ≥1.72.0 recommended
enum34	1.1.10	No	–	–	Yes	Deprecated; for Python <3.4 only
browser-use	0.1.41	Unknown	–	–	Unknown	Check maintainer and origin manually
httpx	0.28.1	Yes	CVE-2024-24773	Moderate	No	DNS rebinding protection issue
Jinja2	3.1.6	No	–	–	No	Actively maintained

5 Agent Support Team

- Cover the full security-to-decision pipeline
- Force constructive conflict
- Simulate real-world decision dynamics
- Reduce blind spots
- Give leadership a balanced recommendation

🚀 Advanced Multi-Agent Orchestration System

Enterprise-grade AI orchestration with intelligent agent selection, consensus analysis, and outlier detection

🔍 Query Analysis

⚙️ Agent Management

📊 System Status

🎯 Submit Query for Multi-Agent Analysis

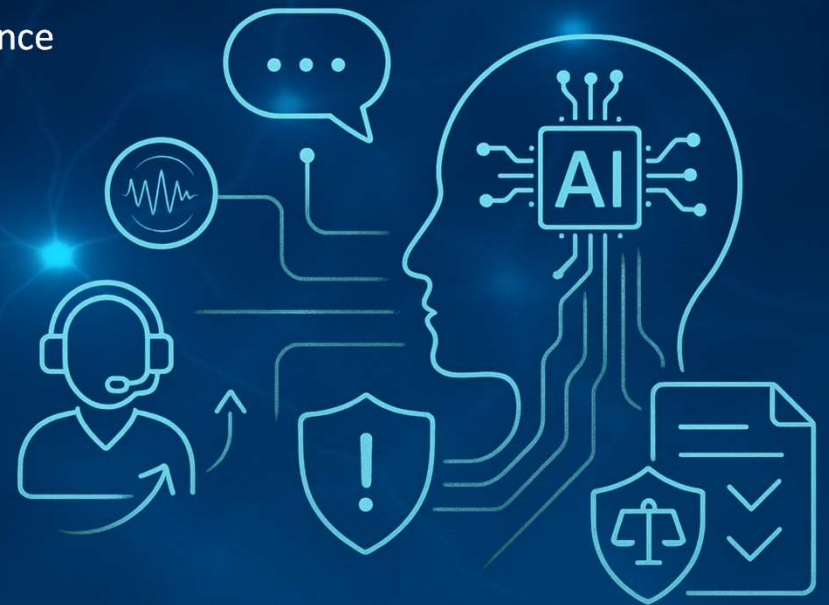
Enter your query for comprehensive multi-agent analysis... (e.g., 'cybersecurity threat analysis', 'physical security assessment for office building', 'financial risk analysis for new product launch')

🚀 Analyze Query

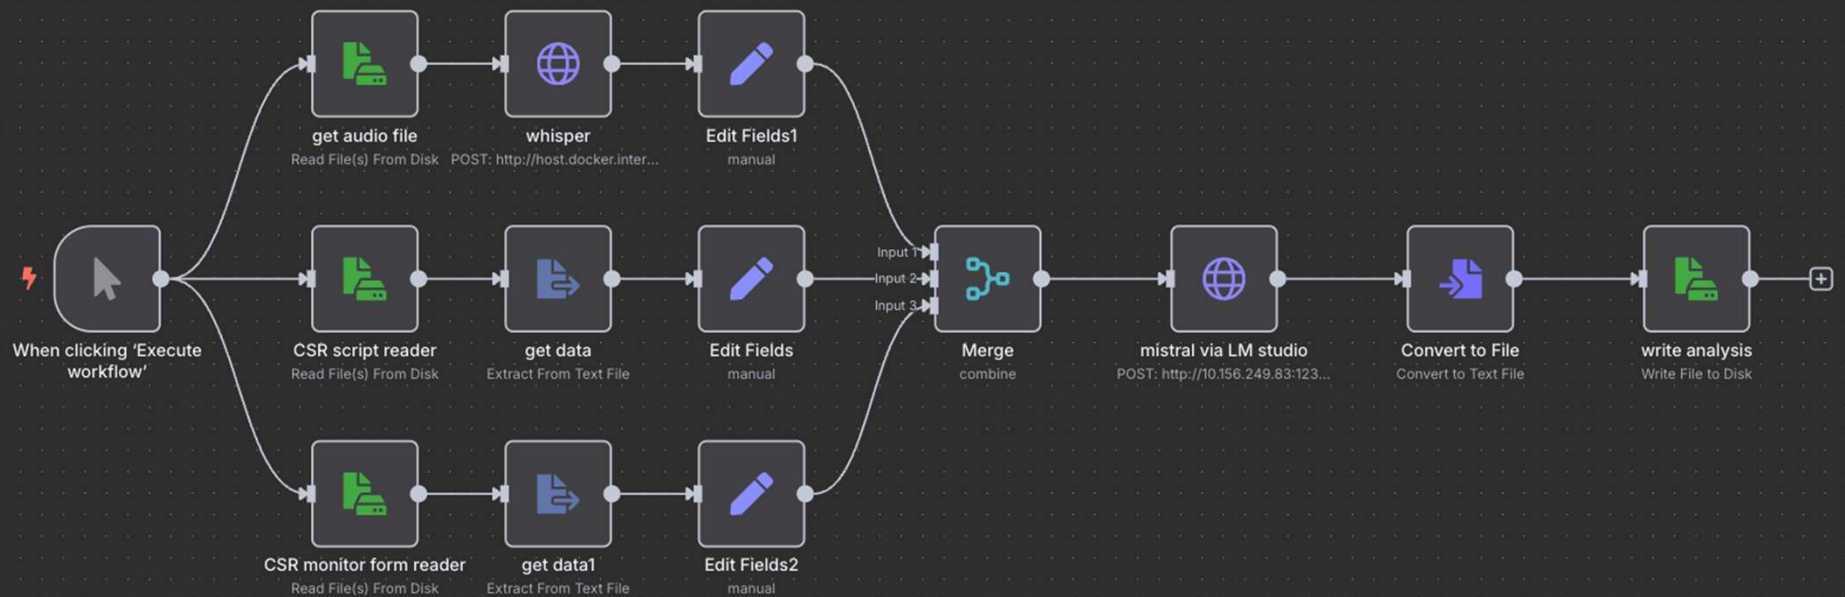
Call Center Analysis

This was a project for our interns which focused on developing a **local, AI-driven proof-of-concept** to analyze customer service interactions using large language models (LLMs). The primary goal was to demonstrate the feasibility of leveraging AI for:

1. Improving customer service representative (CSR) performance
2. Detecting potential fraud or policy deviations
3. Mapping call content against internal policy standards



Call Center Analysis (cont.)



Call Center Analysis

```
77 5. Detection of Potential Fraudulent Activity (You will be given the transcript of the call below):
78 -Use the following questions to assign if there's any potential fraud and malicious intent by the caller. Assign a score and give reasons for it
79 Did the agent follow all required verification steps (DOB, customer ID, address)?
80 Did the caller struggle or hesitate when answering verification questions?
81 Were any answers to identity verification incorrect or inconsistent?
82 Did the caller refuse to provide standard identity verification information?
83 Did the agent bypass any verification steps?
84 Were alternative verification methods used unusually early?
85
86 Did the caller repeat certain requests without answering identity questions?
87 Was the caller overly aggressive, pushy, or evasive?
88 Did the caller attempt to control the conversation or avoid giving direct answers?
89 Did the caller attempt to use urgency or manipulation ("I need this now", "My license is suspended and I'll lose my job")?
90 Did the caller refer to internal DMV knowledge ("I know your system") that a regular customer wouldn't know?
91
92 Did the agent follow the script for sensitive operations (record updates, transfers)?
93 Were sensitive changes (address, name, document status) made too quickly?
94
95 Did the caller ask for unusual or high-risk actions (e.g., unlock account, reissue ID)?
96 Was the caller asking for information about someone else (e.g., "I'm calling for my brother")?
97 Did the caller attempt to provide partial or misleading information?
98 Did the caller use an unusual reason for needing access (e.g., claims of being an attorney or government agent)?
99
```


Call Center Analysis (cont.)

- This person was not able to answer RBQ's
- They made several calls over a few days with different names and customer numbers.

```
[Speaker 1] (0.00s - 15.44s): Hi, this is C [REDACTED] with DMV giving you a call back.
[Speaker 0] (17.58s - 17.98s): Hello.
[Speaker 1] (18.30s - 20.96s): I'm just trying to log in to my DMV account.
[Speaker 1] (22.70s - 25.44s): And I put in my info and then I just set a call with this number.
[Speaker 0] (26.82s - 28.78s): Okay, let me take a look.
[Speaker 1] (30.00s - 31.60s): Uh, I [REDACTED]-H.
[Speaker 1] (32.62s - 33.30s): First name?
[Speaker 1] (34.56s - 36.06s): Logan, L-O-G-A-N.
[Speaker 0] (38.62s - 39.44s): Date of birth?
[Speaker 1] (41.54s - 42.26s): [REDACTED]
[Speaker 1] (46.26s - 48.70s): And what's your customer number?
[Speaker 0] (48.82s - 49.94s): Is your driver's license number?
[Speaker 1] (49.94s - 62.54s): 86 [REDACTED] 78.
[Speaker 0] (62.54s - 65.50s): Thank you.
[Speaker 1] (65.50s - 66.30s): And your current address?
[Speaker 0] (66.30s - 72.24s): That's a tricky one because I move so much.
[Speaker 1] (72.52s - 77.94s): I think my most recent address on there is 818 [REDACTED] Drive.
[Speaker 1] (78.90s - 80.38s): I don't have that one.
[Speaker 0] (81.52s - 84.24s): Do you have...
[Speaker 0] (84.64s - 91.86s): What's the city, state, and zip code for that one?
[Speaker 0] (93.32s - 96.12s): The Fairfax, Virginia 22032.
[Speaker 0] (97.16s - 98.58s): Is that a current address?
[Speaker 0] (100.34s - 102.96s): I currently live in Charlottesville, Virginia.
[Speaker 0] (103.80s - 109.66s): So do we need to get the address updated in the system, or do you want to leave it as is?
[Speaker 1] (111.22s - 114.52s): To be honest, I move so much, it would probably be easiest to leave it as is.
[Speaker 1] (115.02s - 116.36s): Except for one of my cards, so.
[Speaker 1] (117.22s - 117.86s): Not a problem.
[Speaker 1] (118.98s - 120.56s): I'll keep it over here.
[Speaker 1] (130.22s - 136.52s): With the online account, it looks like because it's never been set up before, you actually have to get a PIN number from us.
[Speaker 1] (137.10s - 147.12s): As long as you answer those correctly, I can issue that PIN over the phone, or I'll issue it to your email.
[Speaker 1] (147.80s - 150.80s): And that's what you need to be able to set that online account up.
[Speaker 1] (152.22s - 152.78s): Accounts good.
[Speaker 1] (153.46s - 160.58s): So the first question they have here is, when you last had your driver's license issued, what DMV location did you go into?
[Speaker 0] (160.72s - 161.96s): What city was it in?
[Speaker 0] (164.38s - 165.94s): God, that's a f--- security question.
[Speaker 0] (167.60s - 168.08s): Jesus.
[Speaker 1] (168.32s - 169.64s): Yeah, that computer generated.
[Speaker 1] (170.52s - 177.34s): It was near, is it in Springfield?
[Speaker 1] (179.00s - 179.88s): Is it not Springfield?
[Speaker 1] (182.28s - 182.68s): Jesus.
[Speaker 1] (183.12s - 189.22s): It was the one near, I know what [REDACTED].
[Speaker 0] (194.76s - 196.02s): [REDACTED].
[Speaker 0] (196.84s - 199.74s): It's not a great question, because I don't know if I can answer that.
```


Call Center Analysis (cont.)

****Customer Sentiment****

- The customer initially appears calm and cooperative, providing necessary information when asked. However, there are signs of frustration as the call progresses, particularly when the CSR is unable to locate the record or provide an answer to the customer's inquiries. The customer becomes more insistent and agitated, expressing dissatisfaction with the service.

****Risk Factors****

- Miscommunication: There seems to be a misunderstanding between the CSR and the customer regarding the purpose of the call, as the customer is requesting information about their email address, text number, phone number, lienholder, and record check, but the CSR does not clarify or confirm the specific issues.

- Policy Violations: The CSR did not verify the customer's identity properly before discussing account details and offering solutions, which could potentially compromise sensitive information.

****Key Performance Indicators (KPIs)****

- The issue was not fully resolved by the end of the call, as the CSR was unable to provide answers to the customer's inquiries and had to transfer the call.

- The CSR was courteous at first but became less professional and respectful as the call progressed, failing to maintain a positive tone throughout. The CSR deviated from the script significantly, especially when they were unable to locate the record or provide answers.

- The CSR did not speak clearly or provide accurate information, which contributed to the customer's frustration.

****Wait Time & Flow Analysis****

- There was a long gap (12 seconds) between the end of the customer's question about their lienholder and the CSR's response, which may have caused some confusion for the customer. The overall pacing of the call felt slow due to the CSR's inability to quickly locate the record or provide answers.

****Recommendations****

1. Train CSRs to confirm the purpose of the call with customers and ask clarifying questions to ensure a clear understanding of the issues at hand.
2. Implement stricter identity verification procedures to protect customer information.
3. Improve the efficiency of record checks and provide CSRs with real-time access to customer records.
4. Provide CSRs with additional training on script compliance and proper call handling techniques to maintain a professional and respectful tone throughout calls.
5. Implement a system for tracking and addressing long wait times or delays in call resolution.

****Call Summary****

The call was about various customer service-related inquiries, including email address, text number, phone number, lienholder, and record check. The customer initially presented calm but became increasingly frustrated as the CSR was unable to provide answers or locate their record. The CSR was courteous at first but lacked professionalism and failed to maintain a positive tone throughout the call. The issue was not fully resolved, and the call required a transfer due to the CSR's inability to locate the customer's record.

Tabletop Exercise Creator

Using the existing roles developed and agency knowledge in it's RAG we gave the team the task of creating a tabletop exercise for responding to a hurricane.

To assess the effectiveness of the Virginia DMV's response to the hurricane tabletop exercise by analyzing responses to scenario injects, identifying gaps or concerns, and providing actionable recommendations and discussion points for the upcoming hot wash session.

1. Collect and Organize Responses
2. Assess Each Response Against Objectives
3. Identify Gaps and Areas of Concern
4. Recommend Next Steps
5. Prepare Hot Wash Materials

Deliverables

- A consolidated evaluation report (with findings, gaps, recommendations)
- A scorecard or dashboard of inject performance
- A summary handout or slide deck for the hot wash
- Facilitator notes or talking points for discussion leaders



Tabletop Exercise Creator (cont.)

The AI team created:

- An introduction email to share with the staff members participating
- A complete plan to be executed in 4 phases
- Created all injects based on parameters we established and it's knowledge of the DMV.
- Analyzed all responses to the injects (121 received, 3-7 pages each)
- Created follow-up injects to address a power outage that met our criteria. (used geo-location points from Everbridge to target 15-20% of workforce)
- Produced a detailed executive report with a scorecard, this is feeding back into process improvement efforts.

The screenshot displays the Tabletop Exercise Creator interface. It features a map of Virginia with numerous blue location pins indicating inject points. Overlaid on the map is a text box titled "Key Issues & Conflicts from Week 1 Responses:" which lists several categories of issues identified during the exercise. To the right of the map, there is a grid of document thumbnails, each representing a response or update document, with titles like "Week 3 Responses CSCs and HQ.docx" and "Update Virginia DMV Hurricane Tabletop Exercise Plan-week...".

Key Issues & Conflicts from Week 1 Responses:

- Inconsistent Planning for Mail-Dependent Processes:**
 - Legal Affairs units highlighted that access to HQ for urgent bankruptcy mail is critical and lacks alternatives.
 - No backup plan exists for physical mail processing across the agency, though several units flagged it as a vulnerability.
- Uneven Awareness of Facility Flood Risks:**
 - Several units (e.g., Internal Audit, Data Analytics, PMO) indicated HQ vulnerability but didn't reference basement risk or print shops that others (IT, Weigh Stations) did.
- Limited Coordination Plans for Critical Dependencies:**
 - Many business units rely on IT, Mailroom, and Facilities but fail to specify coordination protocols (e.g., who to contact, how).
 - Example: Hearing Office depends on DMV witnesses and Insurance Services, but coordination protocols are loosely defined.
- Resource & Equipment Protection Gaps:**
 - Only some units (e.g., Weigh Stations, TPSU, IT) discussed securing scale pits, electronics, or relocation of equipment.
 - Others omitted strategies for protecting key infrastructure despite likely flooding (e.g., Legal Affairs, Communications Hq).
- Employee Check-in & Accountability Inconsistency:**
 - Units vary widely in their approach to staff accountability (group text, Everbridge, ad hoc reporting).
 - No central accountability or communication standard is evident.
- Underuse of JOC Support:**
 - Multiple units deferred JOC engagement ("N/A" or "No") despite relying on situational updates, especially regarding facility status.

Areas for Improvement:

- Agency-Wide Mail Access & Virtualization Strategy:**
 - Create a plan for remote access to physical mail (e.g., critical bankruptcy filings), or establish a scanning team to digitize mail.
- Centralized Communication Protocols:**
 - Standardize the use of Everbridge, text chains, or Teams for emergency employee check-ins and status reporting.
- Better Use of Facilities & IT Coordination:**
 - Mandate that all units report what equipment must be protected and involve Facilities/IT in any move or waterproofing efforts.
- Clarify Mutual Aid and Role Flexibility:**
 - Encourage units to identify cross-training or backup assignments to support other teams with skeleton crews (e.g., TPSU, IT, Facilities).
- Resource Tracking During Outages:**
 - Some units (e.g., Weigh Stations, IT) raised generator fuel concerns and scale infrastructure, but few discussed how to track and manage resources during outages.

Next Steps

01

Build a complete AI framework for the agency that weaves AI into the existing fabric of operations, support, security, and training. We need to train our staff and develop future leaders on both current AI applications and emerging capabilities.

02

Complete the Cost Benefit Analysis of moving to dedicated GPU's hosted at QTS for the LLM's and MSI servers to host the n8n instances or leveraging AskSage.AI, or a Hybrid of the two.

03

Continue building our compliance review pipeline.

04

Continue to “make” more time so we can expand on what we've built.

Thank You

"No neural networks were traumatized in the making of this presentation... but a couple may now have a healthy respect for our Security/Risk team."